

CLAIMS

- 1 1. A computer system configured by stored instructions as a garbage collector that
2 reclaims for reuse memory allocated by a mutator executing on the computer system,
3 wherein the garbage collector:
 - 4 A) repeatedly performs global marking operations on the heap;
 - 5 B) treats the heap as divided into a plurality of heap regions for each of at
6 least some of which the garbage collector so maintains a respective re-
7 membered set associated therewith that, independently of which other
8 heap regions contain references to objects in that heap region, entries in
9 that remembered set identify the locations of all such references;
 - 10 C) performs space-incremental-collection operations, associated with respec-
11 tive collection sets, in which the garbage collector employs each remem-
12 bered set associated with a region in the collection set to determine
13 whether objects in the collection set satisfy a potential-reachability crite-
14 rion and reclaims memory occupied by objects that do not satisfy the po-
15 tential-reachability criterion; and
 - 16 D) selects regions for the collection sets by performing calculations of collec-
17 tion-efficiency estimates based at least in part on at least one said global
18 marking operation's results.
- 1 2. A computer system as defined in claim 1 wherein the calculations of the collec-
2 tion-efficiency estimates include, for regions containing only objects allocated at the be-
3 ginning of the most-recent completed global marking operation, estimating from that
4 global marking operation's results respective amounts of memory likely to be reclaimed
5 if those regions are included in the collection set.
- 1 3. A computer system as defined in claim 2 wherein the global marking operations
2 are performed at least in part concurrently with the mutator's execution.

1 4. A computer system as defined in claim 2 wherein the calculations of collection-
2 efficiency estimates include determining which of a set of candidate groups of regions
3 yields the greatest ratio of likely memory reclamation to reclamation cost.

1 5. A computer system as defined in claim 2 wherein:

2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects
5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-
7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and

9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:

- 11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
13 ii) reclaiming the memory space occupied by such objects.

1 6 A computer system as defined in claim 2 wherein the calculations of the collec-
2 tion-efficiency estimates include, for at least some regions, calculating a cost from the
3 sizes of the remembered sets associated therewith.

1 7. A computer system as defined in claim 6 wherein:

2 A) the collector associates respective age values with the regions;

3 B) the calculations of the collection-efficiency estimates include, for some
4 regions, making estimates, based of those regions' age values, of respec-
5 tive amounts of memory likely to be reclaimed if those regions are in-
6 cluded in the remembered set.

- 1 8. A computer system as defined in claim 7 wherein:
2 A) the collector calculates, for each of a plurality of the age values, a respec-
3 tive average of how much memory has been reclaimed from regions with
4 which it associates that age value; and
5 B) the estimates based on the regions' age values are calculated from those
6 averages.

- 1 9. A computer system as defined in claim 1 wherein the calculations of collection-
2 efficiency estimates include determining which of a set of candidate groups of regions
3 yields the greatest ratio of likely memory reclamation to reclamation cost.

- 1 10. A computer system as defined in claim 9 wherein:
2 A) each space-incremental-collection operation includes processing d re-
3 membered-set-log-buffer entries, $d \geq 0$, in order to update remembered
4 sets; and
5 B) the reclamation cost $V(cs)$ for a candidate group cs of regions is deter-
6 mined in accordance with

7
$$V(cs) = V_{fixed} + U \cdot d + \sum_{r \in cs} (S \cdot rsSize(r) + C \cdot liveBytes(r)),$$

8 where V_{fixed} represents fixed costs common to all pauses, d is the number
9 of remembered-set-log-buffer entries to be scanned during that space-
10 incremental-collection operation, U is the average cost of scanning a re-
11 membered-set-log-buffer entry, S is the cost per remembered-set entry of
12 scanning a remembered set, $rsSize(r)$ is the number of remembered-set en-
13 tries in the remembered set maintained for region r , C is the cost per byte
14 of evacuating and scanning an object that is not reclaimed, and $live-$
15 $Bytes(r)$ is an estimate of how many bytes will not be reclaimed from re-
16 gion r .

- 1 11. A computer system as defined in claim 1 wherein:
- 2 A) each of at least some of the global marking operations includes tracing ref-
- 3 erence chains from a root set and so making marks associated with that
- 4 global marking operation and with the locations of respective objects
- 5 thereby encountered that in at least some portions of the heap an object's
- 6 lack of reachability can be inferred at the end of that global marking op-
- 7 eration from the absence of a mark associated with that object's location
- 8 and that global marking operation; and
- 9 B) each of at least some of the space-incremental-collection operations in-
- 10 cludes:
- 11 i) inferring from the marks made by one of the global marking opera-
- 12 tions that objects in the collection set are unreachable; and
- 13 ii) reclaiming the memory space occupied by such objects.
- 1 12. A computer system as defined in claim 11 wherein the calculations of collection-
- 2 efficiency estimates include determining which of a set of candidate groups of regions
- 3 yields the greatest ratio of likely memory reclamation to reclamation cost.
- 1 13. A computer system as defined in claim 1 wherein the global marking operations
- 2 are performed at least in part concurrently with the mutator's execution.
- 1 14. A method of employing a computer system to reclaim for reuse memory dynami-
- 2 cally allocated from a heap in the computer system's memory by a mutator executing on
- 3 that computer system, a method comprising:
- 4 A) repeatedly performing global marking operations on the heap;
- 5 B) treating the heap as divided into a plurality of heap regions and, for each
- 6 of at least some of those heap regions, so maintaining an associated re-
- 7 membered set that, independently of which other heap regions contain ref-

- 8 erences to objects in that heap region, entries in that remembered set iden-
9 tify the locations of all such references;
- 10 C) performing space-incremental-collection operations, associated with re-
11 spective collection sets, in which each remembered set associated with a
12 region in the collection set is employed to determine whether objects in
13 the collection set satisfy a potential-reachability criterion and reclaims
14 memory occupied by objects that do not satisfy the potential-reachability
15 criterion; and
- 16 D) selecting regions for the collection sets by performing calculations of col-
17 lection-efficiency estimates based at least in part on at least one said
18 global marking operation's results.
- 1 15. A method as defined in claim 14 wherein the calculations of the collection-
2 efficiency estimates include, for regions containing only objects allocated at the begin-
3 ning of the most-recent completed global marking operation, estimating from that global
4 marking operation's results respective amounts of memory likely to be reclaimed if those
5 regions are included in the collection set.
- 1 16. A method as defined in claim 15 wherein the global marking operations are per-
2 formed at least in part concurrently with the mutator's execution.
- 1 17. A method as defined in claim 15 wherein the calculations of collection-efficiency
2 estimates include determining which of a set of candidate groups of regions yields the
3 greatest ratio of likely memory reclamation to reclamation cost.
- 1 18. A method as defined in claim 15 wherein:
2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects

- 5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-
7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and
9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:
11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
13 ii) reclaiming the memory space occupied by such objects.

1 19 A method as defined in claim 15 wherein the calculations of the collection-
2 efficiency estimates include, for at least some regions, calculating a cost from the sizes of
3 the remembered sets associated therewith.

- 1 20. A method as defined in claim 19 wherein:
2 A) the method further includes associating respective age values with the re-
3 gions;
4 B) the calculations of the collection-efficiency estimates include, for some
5 regions, making estimates, based of those regions' age values, of respec-
6 tive amounts of memory likely to be reclaimed if those regions are in-
7 cluded in the remembered set.

- 1 21. A method as defined in claim 20 wherein:
2 A) the method further includes calculating, for each of a plurality of the age
3 values, a respective average of how much memory has been reclaimed
4 from regions with which that age value is associated; and
5 B) the estimates based on the regions' age values are calculated from those
6 averages.

1 22. A method as defined in claim 14 wherein the calculations of collection-efficiency
2 estimates include determining which of a set of candidate groups of regions yields the
3 greatest ratio of likely memory reclamation to reclamation cost.

1 23. A method as defined in claim 22 wherein:

- 2 A) each space-incremental-collection operation includes processing d re-
3 membered-set-log-buffer entries, $d \geq 0$, in order to update remembered
4 sets; and
5 B) the reclamation cost $V(cs)$ for a candidate group cs of regions is deter-
6 mined in accordance with

7
$$V(cs) = V_{fixed} + U \cdot d + \sum_{r \in cs} (S \cdot rsSize(r) + C \cdot liveBytes(r)),$$

8 where V_{fixed} represents fixed costs common to all pauses, d is the number
9 of remembered-set-log-buffer entries to be scanned during that space-
10 incremental-collection operation, U is the average cost of scanning a re-
11 membered-set-log-buffer entry, S is the cost per remembered-set entry of
12 scanning a remembered set, $rsSize(r)$ is the number of remembered-set en-
13 tries in the remembered set maintained for region r , C is the cost per byte
14 of evacuating and scanning an object that is not reclaimed, and $live-$
15 $Bytes(r)$ is an estimate of how many bytes will not be reclaimed from re-
16 gion r .

1 24. A method as defined in claim 14 wherein:

- 2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects
5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-

- 7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and
9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:
11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
13 ii) reclaiming the memory space occupied by such objects.

1 25. A method as defined in claim 24 wherein the calculations of collection-efficiency
2 estimates include determining which of a set of candidate groups of regions yields the
3 greatest ratio of likely memory reclamation to reclamation cost.

1 26. A method as defined in claim 14 wherein the global marking operations are per-
2 formed at least in part concurrently with the mutator's execution.

1 27. A storage medium containing instructions readable by a computer system to con-
2 figure the computer system as a garbage collector that reclaims for reuse memory allo-
3 cated by a mutator executing on the computer system, wherein the garbage collector:
4 A) repeatedly performs global marking operations on the heap;
5 B) treats the heap as divided into a plurality of heap regions for each of at
6 least some of which the garbage collector so maintains a respective re-
7 membered set associated therewith that, independently of which other
8 heap regions contain references to objects in that heap region, entries in
9 that remembered set identify the locations of all such references;
10 C) performs space-incremental-collection operations, associated with respec-
11 tive collection sets, in which the garbage collector employs each remem-
12 bered set associated with a region in the collection set to determine
13 whether objects in the collection set satisfy a potential-reachability crite-

14 rion and reclaims memory occupied by objects that do not satisfy the po-
15 tential-reachability criterion; and
16 D) selects regions for the collection sets by performing calculations of collec-
17 tion-efficiency estimates based at least in part on at least one said global
18 marking operation's results.

1 28. A storage medium as defined in claim 29 wherein the calculations of the collec-
2 tion-efficiency estimates include, for regions containing only objects allocated at the be-
3 ginning of the most-recent completed global marking operation, estimating from that
4 global marking operation's results respective amounts of memory likely to be reclaimed
5 if those regions are included in the collection set.

1 29. A storage medium as defined in claim 28 wherein the global marking operations
2 are performed at least in part concurrently with the mutator's execution.

1 30. A storage medium as defined in claim 28 wherein the calculations of collection-
2 efficiency estimates include determining which of a set of candidate groups of regions
3 yields the greatest ratio of likely memory reclamation to reclamation cost.

1 31. A storage medium as defined in claim 28 wherein:

2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects
5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-
7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and

- 9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:
11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
13 ii) reclaiming the memory space occupied by such objects.

1 32 A storage medium as defined in claim 28 wherein the calculations of the collec-
2 tion-efficiency estimates include, for at least some regions, calculating a cost from the
3 sizes of the remembered sets associated therewith.

- 1 33. A storage medium as defined in claim 32 wherein:
2 A) the collector associates respective age values with the regions;
3 B) the calculations of the collection-efficiency estimates include, for some
4 regions, making estimates, based of those regions' age values, of respec-
5 tive amounts of memory likely to be reclaimed if those regions are in-
6 cluded in the remembered set.

- 1 34. A storage medium as defined in claim 33 wherein:
2 A) the collector calculates, for each of a plurality of the age values, a respec-
3 tive average of how much memory has been reclaimed from regions with
4 which it associates that age value; and
5 B) the estimates based on the regions' age values are calculated from those
6 averages.

1 35. A storage medium as defined in claim 27 wherein the calculations of collection-
2 efficiency estimates include determining which of a set of candidate groups of regions
3 yields the greatest ratio of likely memory reclamation to reclamation cost.

1 36. A storage medium as defined in claim 35 wherein:

20/573225.8

- 2 A) each space-incremental-collection operation includes processing d re-
3 membered-set-log-buffer entries, $d \geq 0$, in order to update remembered
4 sets; and
- 5 B) the reclamation cost $V(cs)$ for a candidate group cs of regions is deter-
6 mined in accordance with

$$V(cs) = V_{fixed} + U \cdot d + \sum_{r \in cs} (S \cdot rsSize(r) + C \cdot liveBytes(r)),$$

8 where V_{fixed} represents fixed costs common to all pauses, d is the number
9 of remembered-set-log-buffer entries to be scanned during that space-
10 incremental-collection operation, U is the average cost of scanning a re-
11 membered-set-log-buffer entry, S is the cost per remembered-set entry of
12 scanning a remembered set, $rsSize(r)$ is the number of remembered-set en-
13 tries in the remembered set maintained for region r , C is the cost per byte
14 of evacuating and scanning an object that is not reclaimed, and $live-$
15 $Bytes(r)$ is an estimate of how many bytes will not be reclaimed from re-
16 gion r .

1 37. A storage medium as defined in claim 27 wherein:

- 2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects
5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-
7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and
- 9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:
- 11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
13 ii) reclaiming the memory space occupied by such objects.

1 38. A storage medium as defined in claim 37 wherein the calculations of collection-
2 efficiency estimates include determining which of a set of candidate groups of regions
3 yields the greatest ratio of likely memory reclamation to reclamation cost.

1 39. A storage medium as defined in claim 27 wherein the global marking operations
2 are performed at least in part concurrently with the mutator's execution.

1 40. An electromagnetic signal representing instructions readable by a computer sys-
2 tem to configure the computer system as a garbage collector that reclaims for reuse mem-
3 ory allocated by a mutator executing on the computer system, wherein the garbage collec-
4 tor:

- 5 A) repeatedly performs global marking operations on the heap;
- 6 B) treats the heap as divided into a plurality of heap regions for each of at
7 least some of which the garbage collector so maintains a respective re-
8 membered set associated therewith that, independently of which other
9 heap regions contain references to objects in that heap region, entries in
10 that remembered set identify the locations of all such references;
- 11 C) performs space-incremental-collection operations, associated with respec-
12 tive collection sets, in which the garbage collector employs each remem-
13 bered set associated with a region in the collection set to determine
14 whether objects in the collection set satisfy a potential-reachability crite-
15 rion and reclaims memory occupied by objects that do not satisfy the po-
16 tential-reachability criterion; and
- 17 D) selects regions for the collection sets by performing calculations of collec-
18 tion-efficiency estimates based at least in part on at least one said global
19 marking operation's results.

1 41. An electromagnetic signal as defined in claim 42 wherein the calculations of the
2 collection-efficiency estimates include, for regions containing only objects allocated at
3 the beginning of the most-recent completed global marking operation, estimating from
4 that global marking operation's results respective amounts of memory likely to be re-
5 claimed if those regions are included in the collection set.

1 42. An electromagnetic signal as defined in claim 41 wherein the global marking op-
2 erations are performed at least in part concurrently with the mutator's execution.

1 43. An electromagnetic signal as defined in claim 41 wherein the calculations of col-
2 lection-efficiency estimates include determining which of a set of candidate groups of
3 regions yields the greatest ratio of likely memory reclamation to reclamation cost.

1 44. An electromagnetic signal as defined in claim 41 wherein:

2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects
5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-
7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and

9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:

- 11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
13 ii) reclaiming the memory space occupied by such objects.

1 45 An electromagnetic signal as defined in claim 41 wherein the calculations of the
2 collection-efficiency estimates include, for at least some regions, calculating a cost from
3 the sizes of the remembered sets associated therewith.

1 46. An electromagnetic signal as define in claim 45 wherein:

- 2 A) the collector associates respective age values with the regions;
3 B) the calculations of the collection-efficiency estimates include, for some
4 regions, making estimates, based of those regions' age values, of respec-
5 tive amounts of memory likely to be reclaimed if those regions are in-
6 cluded in the remembered set.

1 47. An electromagnetic signal as define in claim 46 wherein:

- 2 A) the collector calculates, for each of a plurality of the age values, a respec-
3 tive average of how much memory has been reclaimed from regions with
4 which it associates that age value; and
5 B) the estimates based on the regions' age values are calculated from those
6 averages.

1 48. An electromagnetic signal as defined in claim 40 wherein the calculations of col-
2 lection-efficiency estimates include determining which of a set of candidate groups of
3 regions yields the greatest ratio of likely memory reclamation to reclamation cost.

1 49. An electromagnetic signal as defined in claim 48 wherein:

- 2 A) each space-incremental-collection operation includes processing d re-
3 membered-set-log-buffer entries, $d \geq 0$, in order to update remembered
4 sets; and
5 B) the reclamation cost $V(cs)$ for a candidate group cs of regions is deter-
6 mined in accordance with

$$V(cs) = V_{fixed} + U \cdot d + \sum_{recs} (S \cdot rsSize(r) + C \cdot liveBytes(r)),$$

where V_{fixed} represents fixed costs common to all pauses, d is the number of remembered-set-log-buffer entries to be scanned during that space-incremental-collection operation, U is the average cost of scanning a remembered-set-log-buffer entry, S is the cost per remembered-set entry of scanning a remembered set, $rsSize(r)$ is the number of remembered-set entries in the remembered set maintained for region r , C is the cost per byte of evacuating and scanning an object that is not reclaimed, and $liveBytes(r)$ is an estimate of how many bytes will not be reclaimed from region r .

- 1 50. An electromagnetic signal as defined in claim 40 wherein:
 - 2 A) each of at least some of the global marking operations includes tracing ref-
3 erence chains from a root set and so making marks associated with that
4 global marking operation and with the locations of respective objects
5 thereby encountered that in at least some portions of the heap an object's
6 lack of reachability can be inferred at the end of that global marking op-
7 eration from the absence of a mark associated with that object's location
8 and that global marking operation; and
 - 9 B) each of at least some of the space-incremental-collection operations in-
10 cludes:
 - 11 i) inferring from the marks made by one of the global marking opera-
12 tions that objects in the collection set are unreachable; and
 - 13 ii) reclaiming the memory space occupied by such objects.

- 1 51. An electromagnetic signal as defined in claim 50 wherein the calculations of col-
2 lection-efficiency estimates include determining which of a set of candidate groups of
3 regions yields the greatest ratio of likely memory reclamation to reclamation cost.

1 52. An electromagnetic signal as defined in claim 40 wherein the global marking op-
2 erations are performed at least in part concurrently with the mutator's execution.

1 53. For employing a computer system to reclaim for reuse memory dynamically allo-
2 cated from a heap in the computer system's memory by a mutator executing on that com-
3 puter system, a garbage collector comprising:

- 4 A) means for repeatedly performing global marking operations on the heap;
- 5 B) means for treating the heap as divided into a plurality of heap regions and,
6 for each of at least some of those heap regions, so maintaining an associ-
7 ated remembered set that, independently of which other heap regions con-
8 tain references to objects in that heap region, entries in that remembered
9 set identify the locations of all such references;
- 10 C) means for performing space-incremental-collection operations, associated
11 with respective collection sets, in which each remembered set associated
12 with a region in the collection set is employed to determine whether ob-
13 jects in the collection set satisfy a potential-reachability criterion and re-
14 claims memory occupied by objects that do not satisfy the potential-
15 reachability criterion; and
- 16 D) means for selecting regions for the collection sets by performing calcula-
17 tions of collection-efficiency estimates based at least in part on at least one
18 said global marking operation's results.